Aging of the Baby Boomers: Demographics and Propagation of Tax Shocks^{*}

Domenico Ferraro

Giuseppe Fiori

Arizona State University

North Carolina State University

November 8, 2015

Abstract

We investigate the consequences of demographic change for the propagation of exogenous tax changes in the U.S. labor market. We document that the responsiveness of unemployment rates to tax changes largely varies across age groups: the unemployment rate response of the young is nearly twice as large as that of prime-age workers. Such heterogeneity is the channel through which shifts in the age composition of the labor force impact the responsiveness of the aggregate U.S. unemployment rate to tax changes. We then present a model with frictional unemployment and learning about match quality that quantitatively accounts for the estimated aggregate unemployment semi-elasticity to tax cuts. We use the calibrated model to quantify the impact of an aging labor force on the propagation of tax shocks. The results indicate that the aging of the Baby Boomers reduces the aggregate unemployment semi-elasticity to tax cuts by approximately 50 percent.

JEL Classification: E24; E62; J11

Keywords: Fiscal policy; Tax changes; Demographics; Unemployment

^{*}Ferraro: Department of Economics, W. P. Carey School of Business, Arizona State University, PO Box 879801, Tempe, AZ 85287-9801 (e-mail: domenico.ferraro@asu.edu); Fiori: Department of Economics, Poole College of Management, North Carolina State University, PO Box 8110, Raleigh, NC 27695-8110 (e-mail: gfiiori@ncsu.edu). First version: November 8, 2015. We thank Alexander Bick, Nir Jaimovich, Valerie Ramey, and Pedro Silos for comments and suggestions. Any errors are our own.

1 Introduction

In this paper, we investigate the consequences of demographic change for the propagation of exogenous tax changes in the U.S. labor market. We ask the question: How do shifts in the age composition of the labor force affect the response of the aggregate unemployment rate to unanticipated tax cuts? The United States represent an ideal laboratory to answer this question as the aging of the Baby Boom provides us with a real-time natural quasi-experiment. We argue that the demographic composition of the labor force constitutes an important propagation mechanism of tax shocks.¹

We document that the responsiveness of unemployment rates to tax shocks largely varies across age groups: the unemployment rate response of the young is nearly twice as large as that of prime-age workers. Documenting these age-specific differences in the responsiveness to tax shocks is the first contribution of the paper. This heterogeneity is the channel through which shifts in the age composition of the U.S. labor force affect the response of the aggregate unemployment rate to tax changes. We then propose a theory to *quantify* the impact of an aging labor force on the propagation of tax shocks. Quantifying these effects is the second contribution of the paper. The results indicate that the aging of the Baby Boom decreases the aggregate unemployment semi-elasticity to tax cuts by nearly 50 percent.

Recently, a great deal of attention has been devoted to studying the role of fiscal policy in stimulating economic activity. The renewed interest has been spurred by the Great Recession of 2007-2009 and the uncertainty surrounding the beneficial effects of the American Recovery and Reinvestment Act of 2009 (ARRA)—the stimulus package enacted to preserve and create jobs, and promote economic recovery.

Not surprisingly, then, several empirical studies have examined the effects of fiscal policy shocks.² No one, however, has yet systematically studied how the aging of the U.S. labor force affects the propagation of tax shocks. We do that in this paper using narrative identification of exogenous tax changes and quantitative theory. We establish the relationship between demographics and the propagation of tax shocks in the following manner.

¹Shimer (1999) shows that the entry of the baby boom into the labor force in the late-1970s and their aging accounts for a substantial fraction of the rise and fall in the unemployment rates observed in the past 50 years, whereas Jaimovich and Siu (2009) show that such demographic change accounts for a significant fraction of the decrease in business cycle volatility since the mid-1980s.

²The bulk of the literature considers aggregate macroeconomic variables (see Ramey and Shapiro, 1998; Romer and Romer, 2010; Monacelli et al., 2010; Barro and Redlick, 2011; Ramey, 2011a; Auerbach and Gorodnichenko, 2012; Mertens and Ravn, 2013, 2014; Mertens, 2013; Cloyne, 2013; Ramey and Zubairy, 2014; Barnichon and Matthes, 2015), whereas Anderson et al. (2011) document heterogeneous effects of fiscal policy shocks on consumption, based on income levels and age.

In Section 2, after isolating exogenous changes in average marginal tax rates in structural vector autoregressions (SVARs) using a narrative identification approach (see Mertens and Ravn, 2013), we document that (i) aggregate job-separation and job-finding rates are jointly responsible for the response of the unemployment rate to tax shocks, and that (ii) the long-run (steady-state) unemployment semi-elasticity is in fact a sufficient statistic to capture the aggregate unemployment response to tax cuts; (iii) age-specific labor force shares do not move in response to tax shocks, whereas (iv) the unemployment rate response of the young is nearly twice as large as that of prime-age workers.

Given these empirical findings, a natural conjecture is that the responsiveness of the aggregate U.S. unemployment rate to tax shocks depends on the age composition of the labor force. When an economy is characterized by a smaller share of young workers, everything else equal, these should be periods of lesser responsiveness to tax cuts. We quantify these effects using a frictional model of the labor market.

In Section 3, we propose a simple theory of frictional unemployment based on learning about occupational fit.³ During the youth, workers switch jobs to learn the occupation in which they are most productive. Hence, young workers are most likely to be in matches of poor occupational fit, whereas prime-age workers are most likely settled in their best occupational fit. As a result, the young generate small rents (wages in excess of the value of unemployment), which makes them extremely sensitive to tax changes. By contrast, prime-age workers generate large rents, which makes them relatively insensitive to such tax changes. We argue that explicitly taking in account these age-specific differences provides new insights into the propagation mechanism of tax shocks.⁴

Quantitatively explaining the *estimated* aggregate unemployment semi-elasticity to tax cuts is the test of the underlying theory. In Section 4, we calibrate the model to post-war U.S. data for 1950-2006 and show that the theory quantitatively explains the observed life-cycle profile of average unemployment rates and that it indeed replicates the estimated semi-elasticity of the aggregate unemployment rate to tax shocks.

In Section 5, we use the theory as a measurement device. Specifically, through the lens of the calibrated model, we reconstruct the U.S. history of unemployment semi-elasticities from 1950 to 2015. The counterfactuals show that the aging of the Baby Boom decreases the aggregate unemployment semi-elasticity by nearly 50 percent. Thus, the theory predicts

³Pries (2004) studies learning as a propagation mechanism of business cycle shocks, whereas Gervais et al. (2014) propose learning as an explanation for the high unemployment rates of the young.

⁴Analogously, Gomme et al. (2005), Hansen and İmrohoroğlu (2009), Dyrda et al. (2012), and Jaimovich et al. (2013) study aggregate implications of age-specific differences in cyclical movements of hours worked.

that the demographic composition of the labor force constitutes a potentially important propagation mechanism of tax shocks. In Section 6, we provide concluding remarks.

2 Labor Market Response to Tax Shocks

In this section, we document new evidence on the dynamic response of the U.S. labor market to narratively identified tax shocks—“narrative tax shocks.” To this aim, we use the “proxy SVAR” approach proposed by [Mertens and Ravn \(2013\)](#). Specifically, we focus on changes in marginal personal income tax rates (AMTR, henceforth), and labor market variables, such as unemployment rate (fraction of unemployed workers in the labor force) and participation rate (labor force as fraction of the population) at the aggregate level and at the disaggregated level by five age groups (20-24, 25-34, 35-44, 45-54, and 55-64). Data on average marginal tax rates has been tabulated by Robert Barro and Charles Redlick. Hence, we refer the reader to [Barro and Redlick \(2011\)](#) for additional details. As in [Mertens and Ravn \(2013\)](#), AMTR consists of two components: the federal individual income tax and the social security payroll tax (FICA).⁵

Narrative tax shocks. Following [Mertens and Ravn \(2013, 2014\)](#), the identification of the true underlying tax shocks is obtained by the means of a SVAR and the use of a proxy for exogenous variation in tax rates, as an external instrument. We consider the time series of “exogenous” changes in tax liabilities constructed by [Romer and Romer \(2010\)](#) as the designed proxy, but we only use observations on tax liability changes legislated and implemented within the year to avoid anticipation effects, as in [Mertens and Ravn \(2012\)](#). These narratively identified, tax liabilities shocks are then used as an instrument to estimate the response of aggregate and age-specific labor market variables to exogenous changes in marginal personal income tax rates in the United States over an annual sample of fifty-seven years, 1950-2006.

Proxy SVAR specification. Our estimates of the dynamic effects of tax changes on the U.S. labor market are based on a VAR with 5 variables:

$$\mathcal{Y}_t = [\text{AMTR}_t, \ln(\text{PITB}_t), \ln(G_t), \mathcal{X}_t, \ln(\text{DEBT}_t)], \quad (1)$$

where (i) AMTR_t is the average marginal personal income tax rate; (ii) PITB_t is the

⁵As in [Barro and Redlick \(2011\)](#), AMTR is the average marginal income tax rate weighted by a concept of income that is close to labor income: wages, self-employment, partnership income, and S-corporation income.

personal income tax base in real per capita terms; (iii) G_t is government purchases of final goods in real per capita terms; (iv) $\mathcal{X}_t \in \{\text{UR}_t, \text{PR}_t, \text{JSR}_t, \text{JFR}_t\}$, where UR_t and PR_t are aggregate unemployment and participation rates, respectively, and JSR_t and JFR_t are job-separation and job-finding rates, respectively, as constructed by Robert Shimer (see [Shimer, 2012](#), for details); and (v) DEBT_t is federal debt in real per capita terms. The sample consists of annual observations for the period 1950-2006. The lag length in the VAR is set to two. All impulse responses are for a 1 percentage point cut in the tax rate and we show results for a forecast horizon of 5 years. Government debt is an important variable to include in the VAR specification: given the government budget constraint, changes in taxes must eventually lead to adjustments in other fiscal instruments, as such we deem appropriate to explicitly allow for the potential feedback from debt to taxes and spending.⁶

2.1 Aggregate results

We now turn to study the dynamic response of the U.S. labor market to tax shocks.⁷ Specifically, we focus on the extensive margin of the labor market; that is, aggregate U.S. unemployment and labor force participation rates. In doing so, we shed light on whether the unemployment or participation margin is the key driving force of the labor market response to narrative tax shocks. To interpret the main empirical findings of this section, we will make use of the following decomposition of the employment to population ratio:

$$\frac{\text{employment}}{\text{population}} = \underbrace{\left(1 - \frac{\text{unemployment}}{\text{employment} + \text{unemployment}}\right)}_{\text{1 minus the unemployment rate}} \times \underbrace{\left(\frac{\text{employment} + \text{unemployment}}{\text{population}}\right)}_{\text{participation rate}}.$$

Figure 1 shows the dynamic response of the unemployment rate to a 1 percentage point cut in AMTR. The peak response is large in magnitude: a 1 percentage point cut in AMTR leads to an approximately 0.7 percentage points decrease in the aggregate U.S. unemployment rate. Note that historically a 1 percentage point cut in AMTR is not an unusual event as the standard deviation of AMTR is 4.6 percent for the period 1950-2006. Hence, these results

⁶[Christ \(1968\)](#) and [Sims \(1998\)](#) warn against policy analysis that fails to keep track of the implications of the government budget constraint. [Burnside et al. \(2004\)](#) and [Ramey \(2011b\)](#) argue that the effects of shocks to government purchases may differ depending on the endogenous response of other fiscal instruments.

⁷Figure B.1 in Appendix B replicates the response of the real GDP per capita to a 1 percentage point cut in AMTR, as in [Mertens and Ravn \(2013\)](#). We refer the reader to [Mertens and Ravn \(2014\)](#) for a thorough comparison of the estimated output effects found in the SVAR literature.

point to a highly significant and quantitatively large effect of tax cuts on unemployment.

Figure 2 shows the dynamic response of the participation rate to an equally-sized cut in AMTR. In contrast with the results for the aggregate unemployment rate in Figure 1, the response of the labor force participation rate is statistically and economically insignificant.

These results imply that the overall response of the employment to population ratio is the mirror image of the response of the unemployment rate. And that a theory of unemployment is hence the natural framework to tackle the research question of this paper.

2.2 Decomposing separation- versus hiring-driven effects

Can the Mortensen-Pissarides (MP) search-and-matching model account for the estimated response of the unemployment rate to narrative tax shocks? And what are the propagation mechanisms of such tax shocks? To answer this question, we next focus on the reduced-form implications of the typical MP model (see Pissarides, 1985; Mortensen and Pissarides, 1994). In such a model, changes in unemployment rates are driven by the fraction of employed workers who lost jobs in the previous period minus the unemployed workers who found jobs:

$$u_{t+1} - u_t = s_t e_t - f_t u_t, \quad (2)$$

where the variables u_t and $e_t = 1 - u_t$ are unemployment and employment stocks, and s_t and f_t are job-separation and job-finding rates, respectively. Thus, equation (2) identifies the job-separation rate (“separation margin”) and the job-finding rate (“hiring margin”) as the key driving forces of changes in the aggregate unemployment rate.

Hence, we now turn to discuss the estimated dynamic responses to narrative tax shocks of the two flow rates: s_t and f_t . Based on publicly available data on unemployment and employment from the Current Population Survey (CPS), time series for s_t and f_t are easily constructed (see Shimer, 2012).⁸ Figure 3 shows the responses of such constructed series to a 1 percentage point cut in AMTR. In panel A, the job-separation rate drops by -0.14 percentage points at the trough of the response, whereas, in panel B, the job-finding rate raises by approximately 3 percentage points at the peak of the response. We conclude, then, that both responses are qualitatively consistent with the unemployment rate response in Figure 1, and in accord with the equilibrium reduced-form of the typical MP model in (2).

⁸Job-separation and job-finding rates are constructed under two assumptions: (1) workers do not transit in and out of the labor force; and (2) workers are homogeneous with respect to job-finding and job-separation probabilities. I refer the reader to Shimer (2012) for further details.

However, to quantify the relative contribution of separation and hiring margin, we need to derive a decomposition of the unemployment rate response to tax shocks that compares the contribution of job-separation and job-finding rates on an equal footing. To this aim, we proceed in two steps.

First, we rely on a theoretical approximation argument, that hinges on the irrelevance of turnover (transitional) dynamics for empirically relevant values of s_t and f_t . In the absence of any type of disturbance, the unemployment rate in (2) converges to the theoretical steady-state value $u_t^{\text{ss}} \equiv \frac{s_t}{s_t + f_t}$, which only depends on contemporaneous values of job-separation and job-finding rates. Hall (2005) and Shimer (2012) show that, at the quarterly frequency, such steady-state approximation generates values that are nearly indistinguishable from the actual unemployment rates observed in U.S. data: that is, $u_t^{\text{ss}} \simeq u_t$ at the quarterly frequency (and thus at lower frequencies as well). This happens because, in the data, $s_t + f_t$ is typically close to 0.5 on a monthly basis, such that the half-life of a deviation from the steady-state unemployment rate is approximately one month. In our sample, monthly job-separation and job-finding rates are 3.4 and 45.3 percent on average, respectively, which suggests that the conditions for the steady-state approximation to work are indeed met; and even more so for the annual frequency our estimates are based on.

Second, based on Elsby et al. (2009), Fujita and Ramey (2009), and Pissarides (2009), we decompose the changes in the (steady-state approximation of) actual unemployment rates into the contribution due to changes in the job-separation rate and the contribution due to changes in the job-finding rate:

$$du_t^{\text{ss}} \approx \bar{u}^{\text{ss}} (1 - \bar{u}^{\text{ss}}) \frac{ds_t}{\bar{s}} - \bar{u}^{\text{ss}} (1 - \bar{u}^{\text{ss}}) \frac{df_t}{\bar{f}}, \quad (3)$$

where $dx_t = x_t - \bar{x}$ indicates deviations of the generic variable x_t from its sample average $\bar{x}$, and $\bar{u}^{\text{ss}} \equiv \bar{s} / (\bar{s} + \bar{f})$, where $\bar{s}$ and $\bar{f}$ are sample averages of job-separation and job-finding rates, respectively. Equation (3) provides an additive decomposition in which the contributions of job-separation and job-finding rates are comparable on an equal footing with respect to their impact on the observed changes in the unemployment rate. Note that equation (3) holds as equality only for infinitesimal changes. For discrete changes, instead, it is only an approximation. However, Fujita and Ramey (2009)'s regression-based estimation of the decomposition in (3) verifies that *unconditionally* it works well for discrete changes.

We next show that such decomposition works well also for the *conditional* response of the U.S. unemployment rate to narrative tax shocks. Specifically, the counterfactual response

implied by equation (3) is

$$\widehat{du}_h^{\text{ss}} = \underbrace{\bar{u}^{\text{ss}} (1 - \bar{u}^{\text{ss}}) \frac{\widehat{ds}_h}{\bar{s}}}_{du_h^{\text{jsr}}} - \underbrace{\bar{u}^{\text{ss}} (1 - \bar{u}^{\text{ss}}) \frac{\widehat{df}_h}{\bar{f}}}_{du_h^{\text{jfr}}}, \quad (4)$$

where h denotes the number of years after the shock and $\widehat{ds}_h$ and $\widehat{df}_h$ are the estimated responses of job-separation and job-finding rates, respectively, as shown in Figure 3.

Figure 4 shows that the counterfactual response of the unemployment rate under the steady-state approximation, $\widehat{du}_h^{\text{ss}}$, (dashed line with diamonds) is remarkably close to the estimated response of the actual unemployment rate (full line with circles). This result shows that, at the annual frequency, turnover dynamics is in fact negligible for understanding the response of the actual unemployment rate to tax shocks. As a result, the *steady-state* semi-elasticities of job-separation and job-finding rates with respect to unanticipated changes in tax rates are sufficient statistics to quantify the response of the unemployment rate to tax shocks. Specifically, the semi-elasticities implied by the estimates for the job-separation and job-finding rate are $\frac{d \ln s_t}{\epsilon_t^{\text{AMTR}}} = 3.29$ and $\frac{d \ln f_t}{\epsilon_t^{\text{AMTR}}} = 2.42$, respectively, whereas the implied semi-elasticity for the unemployment rate is $\frac{d \ln u_t}{\epsilon_t^{\text{AMTR}}} = 5.94$, where ϵ_t^{AMTR} captures the purely *exogenous* changes to average tax rates.

Table 1: Separation- versus hiring-driven tax effects

h (years after shock):	1	2	3	4	5
du_h^{jsr} as % of $\widehat{du}_h^{\text{ss}}$	0.58	0.40	0.36	0.35	0.42
du_h^{jfr} as % of $\widehat{du}_h^{\text{ss}}$	0.42	0.60	0.64	0.65	0.58

Notes: See equation (4) for definitions of $\widehat{du}_h^{\text{ss}}$, du_h^{jsr} , and du_h^{jfr} .

In Table 1, we use the decomposition in equation (4) to quantify the relative contribution of job-separation and job-finding rates to the response of the actual U.S. unemployment rate to narrative tax shocks. The results show that the job-separation rate accounts for approximately 60 percent of the overall unemployment rate response in the first year after the shock, whereas the split between job-separation and job-finding rates is inverted from the second year after the shock onward. These results indicate that to fully account for the effects of tax shocks, a theory must rely on an active separation margin.

2.3 Age-specific results

We now turn to the demographics of the U.S. labor market response to tax shocks. Specifically, let us consider the following decomposition of the aggregate unemployment rate:

$$\frac{\text{unemployment}}{\text{labor force}} = \sum_{a \in \mathcal{A}} \underbrace{\frac{\text{labor force}_a}{\text{labor force}}}_{\text{age-specific labor force share}} \times \underbrace{\frac{\text{unemployment}_a}{\text{labor force}_a}}_{\text{age-specific unemployment rate}}, \quad (5)$$

where $\mathcal{A} = \{16-19, 20-24, 25-34, 35-44, 45-54, 55-64, 65+\}$ is a list of age groups, whereas the subscript “ a ” indicates an age group in that list; the labor force is the sum of employment and unemployment, as defined by the Bureau of Labor Statistics (BLS). The decomposition in (5) suggests that the estimated response of the aggregate unemployment rate to narrative tax shocks may be due to either the response of age-specific labor force shares or the response of age-specific unemployment rates, or both.

Next, to disentangle the contribution of labor force shares from that of the unemployment rates, we construct a counterfactual series of the aggregate unemployment rate in which age-specific labor force shares are fixed at their sample averages, $\overline{\text{LFS}}_a$, for 1950-2006, whereas we let age-specific unemployment rates, $\text{UR}_{a,t}$, vary over time:

$$\text{UR}_t^{\text{CTRFL}} = \sum_{a \in \mathcal{A}} \overline{\text{LFS}}_a \times \text{UR}_{a,t}. \quad (6)$$

Then, we estimate the proxy SVAR by replacing the actual unemployment rate with the counterfactual series with fixed labor force shares in (6). In Figure 4, panel B shows that the response to the tax shock of the counterfactual series (dashed line with diamonds) is indeed remarkably close to the response of the actual unemployment rate (full line with circles), which suggests that the key drivers of the aggregate unemployment rate response to narrative tax shocks are the age-specific unemployment rates. Figure 5 shows that such result holds for the 20-64 age group as well, that is the age group we focus on in the theory.

These results are important for the scope of this paper as they validate the exclusion restriction that justifies our approach. Age-specific labor force shares depend on demographic trends underlying the U.S. labor market, that are in fact unaffected by the tax shocks we consider. As such, we can study how such demographic trends, and so changes in the age composition of the labor force, impact the propagation of tax shocks.

Figure 6 shows the responses of age-specific unemployment rates to a 1 percentage point cut in AMTR. The estimates display stark differences in the responses of young and prime-age workers. In panel A, the unemployment rate for the 20-24 years old falls by 1 percentage point at the peak of the response, that is approximately two times as large as the peak response for the 35-64 years old.⁹ This age-specific heterogeneity in the responsiveness to tax shocks is the channel through which shifts in the age composition of the U.S. labor force affect the response of the aggregate unemployment rate to exogenous changes in taxes.

Figure 1: Unemployment Rate Response to a Tax Cut

Notes: The figure shows the response to a 1 percentage point cut in the average marginal personal income tax rate. Full lines with circles are point estimates; dashed lines are 95 percent confidence bands.

⁹Figure B.2 in Appendix B shows the responses of age-specific participation rates to a 1 percentage point cut in AMTR. The responses of the participation rate for the young and prime-aged workers are not statistically and economically significant.

Figure 2: Participation Rate Response to a Tax Cut

Notes: The figure shows the response to a 1 percentage point cut in the average marginal personal income tax rate. Full lines with circles are point estimates; dashed lines are 95 percent confidence bands.

Figure 3: Job-Separation and Job-Finding Rate Response to a Tax Cut

Notes: The figure shows the response to a 1 percentage point cut in the average marginal personal income tax rate. Full lines with circles are point estimates; dashed lines are 95 percent confidence bands. Data for job-separation and job-finding rates was constructed by Robert Shimer (see [Shimer, 2012](#), for details).

Figure 4: Counterfactual Unemployment Rate Response to a Tax Cut

Notes: In panel A and B, full lines with circles are point estimates for the response to a 1 percentage point cut in the average marginal personal income tax rate; dashed lines are 95 percent confidence bands. In panel A, dashed line with diamonds shows the steady-state approximation of the response for the actual unemployment rate, as implied by equation (4). In panel B, dashed line with diamonds shows the response of the counterfactual unemployment rate with fixed age-specific labor force shares.

Figure 5: Unemployment Rate Response to a Tax Cut for the 20-64 Age Group

Notes: Full lines with circles are point estimates for the response to a 1 percentage point cut in the average marginal personal income tax rate; dashed lines are 95 percent confidence bands. Dashed line with diamonds shows the response of the counterfactual unemployment rate with fixed age-specific labor force shares.

Figure 6: Age-specific Unemployment Rate Response to a Tax Cut

Notes: The figure shows the response to a 1 percentage point cut in the average marginal personal income tax rate. Full lines with circles are point estimates; dashed lines are 95 percent confidence bands.

3 The Model

In this section, we present the model used to analyze the quantitative impact of tax changes on the U.S. unemployment rate and the consequences of the aging of the U.S. labor force for the propagation of tax shocks. The key ingredients of the model are a frictional labor market and learning about occupational fit. The model builds on [Pries \(2004\)](#) and [Gervais et al. \(2014\)](#). Specifically, [Pries \(2004\)](#) studies learning about match quality as a propagation mechanism of business cycle shocks, whereas [Gervais et al. \(2014\)](#) propose learning about occupational fit as an explanation for the high job-separation rates and unemployment rates of the young. We propose learning as a propagation mechanism of tax shocks.

3.1 Environment

Time is discrete and continues forever indexed by $t \geq 0$. Throughout, we omit time subscripts unless needed for clarity.

Agents. There is a measure 1 of ex-ante identical workers either employed, $e \in [0, 1]$, or unemployed, $u \in [0, 1]$, and searching for a job. Each worker is endowed with an indivisible unit of time, that can be either supplied as labor services when employed or used for search activities when unemployed. We adopt a 2-state representation of the labor market, as such we abstract from movements in and out of the labor force: $e + u = 1$ for all $t \geq 0$.

We consider a model economy consisting of M occupations, where $M \geq 2$ is an integer. Specifically, the M occupations are identical, except for the “name” $m \in \{1, \dots, M\}$. Each worker is best-suited for one (and one only) occupation, $m^* \in \{1, \dots, M\}$, which we think of the worker’s best occupational fit (“true calling”). Workers are infinitely lived and die with probability $1 - \sigma$, with $\sigma \in (0, 1)$. To preserve stationarity, we assume that in each period a mass of new workers enters the labor force replacing the workers that have died. New entrants do not know their true calling, but they know that m^* is randomly assigned into one of the M occupations with probability $1/M$. Such assignment is independently distributed across all new entrants. In order to learn her best occupational fit, the worker must search, be matched, and work in that occupation. Due to the presence of learning frictions, it takes time for the worker to learn her true calling. Specifically, we assume learning occurs at an exogenous and constant rate $\lambda \in (0, 1)$. Note that workers are ex-ante identical, yet ex-post heterogeneity in labor market histories arises in equilibrium due to stochastic matching, separations, and learning. Worker heterogeneity is summarized by the worker’s

type, $i \in \{1, \dots, M\}$: the index i counts the number of ill-suited occupations sampled by the worker, plus one. Workers know that their true calling is uniformly distributed over all M occupations. Hence, an unemployed worker of type $i = 1$ has a flat prior over the best occupational fit of $\pi_1 = 1/M$, as such she randomly selects an occupational market $m \in \{1, \dots, M\}$ to search for a job. If occupation m turns out to be not her true calling, the worker becomes of type $i = 2$, as such she understands that her true calling is now uniformly distributed over the remaining $M - 1$ occupations. Specifically, a worker of type $i < M$ has previously worked at $i - 1$ ill-suited occupations and therefore she has $M - i + 1$ occupations left to try. Therefore, the probability that a worker of type i has found her true calling is $\pi_i = 1/(M - i + 1)$. A worker of type $i = M$ knows her best occupational fit with certainty either because she has tried $M - 1$ ill-suited occupations or because she has been “lucky.”

A typical worker can then be in three different states: (1) working in her true calling, and producing $y_G + x$ units of output—“good” occupational match; (2) working in an ill-suited occupation, and producing $y_B + x$, with $y_B < y_G$ —“bad” occupational match; and (3) in the learning phase, and producing $\bar{y}_i + x$, with $\bar{y}_i = \pi_i y_G + (1 - \pi_i) y_B$, where $y_B < \bar{y}_i \leq y_G$ —“unknown” occupational match. Match quality $y \in \{y_B, \bar{y}_i, y_G\}$ is contaminated by match-specific shocks, x , that are identically and independently distributed across time and across matches, and distributed uniformly on $[-x_u, x_u]$: $x \stackrel{iid}{\sim} F(x)$, where $F(x)$ is the cumulative distribution function (CDF) of a continuous uniform distribution.

The economy is also populated by a continuum of identical employers, either producing output or posting vacancies to hire unemployed workers. The worker’s type i is known by the worker, but it becomes observable to the employer only after posting a vacancy and meeting a worker. In this sense, search and matching are random. As standard in the literature, the mass of employers is endogenously determined in free-entry equilibrium. Workers and employers are infinitely lived, have risk-neutral preferences, and discount payoffs at rate $\beta \in (0, 1)$.

Matching technology. We adopt the standard frictional view of the labor market and postulate the existence of a constant returns to scale (CRS) matching technology that determines the number of matches between employers and workers as function of unemployed workers seeking jobs, u_t , and job vacancies posted by employers, v_t :

$$h(v_t, u_t) = \xi v_t^\alpha u_t^{1-\alpha}, \text{ with } \alpha \in (0, 1),$$

where $h(v_t, u_t)$ denotes worker-employer matches, and the parameter ξ is a shifter that

scales the efficiency of the matching process. Given the CRS assumption, conditions in the labor market are summarized by the labor market tightness $\theta_t \equiv v_t/u_t$. An unemployed worker finds a job with probability $p(\theta_t) = h(v_t, u_t)/u_t = \xi\theta_t^\alpha$, and a job vacancy is filled with probability $q(\theta_t) = h(v_t, u_t)/v_t = p(\theta_t)/\theta_t = \xi\theta_t^{\alpha-1}$. In a tight labor market, it is easy for job seekers to find jobs—the job-finding probability $p(\theta_t)$ is high—and difficult for employers to hire—the job-filling probability $q(\theta_t)$ is low. The cost of keeping a job vacancy open has a per-period cost k . As such, the expected cost of opening and maintaining a job vacancy is $k/q(\theta_t)$.

Nash-bargaining. Upon meeting, employers and workers enter bilateral, generalized Nash-bargaining that determines the wage payment to the worker. The Nash-bargaining solution implies that worker and employer receive a constant and proportional share of the total match surplus, which in turn results in separation decisions that are bilaterally efficient. This result allows us to conveniently and parsimoniously work with the Bellman equations that describe the expected total surplus of bad, good, and unknown matches, rather than with pairs of separate Bellman equations for the worker and employer. With this approach, wage determination is not explicitly treated. Rather it is implicitly represented in the Bellman equations via the worker’s bargaining parameter, τ . In the background, associated with each of the three types of matches, there is a corresponding wage that achieves the Nash-bargaining outcome.

Taxes. We consider a proportional tax on labor income, τ_W .¹⁰ Under Nash-bargaining, if taxes are the only distortion, then it is irrelevant whether taxes are levied on the employer or the worker—all that matters is the tax levied on the match. As a result, the introduction of the proportional tax rate τ_W mandates only minor changes to the surplus equations of bad, good, and unknown matches as it effectively creates a wedge between the worker’s outside option, b , and what we refer to as the *effective* outside option, $z = b/(1 - \tau_W)$. As such, unexpected changes in τ_W are tantamount to shocks to the relative return of market versus non-market activity.

3.2 Equilibrium

We now turn to discuss the stationary equilibrium of the model. Under Nash-bargaining, the equilibrium of this model economy is fully characterized by (a system of) Bellman equations for the total match surplus in good, bad, and unknown occupational matches.

¹⁰One can show that the proportional tax rate τ_W can be interpreted as the average marginal tax rate faced by a representative household that pools the labor income of all the agents in the economy.

Surplus equations. After the realization of the match-specific shock, x , workers and employers, with either match quality, jointly decide whether to remain in the match, and produce output, or to destroy the match. Nonetheless, workers at different stages of their labor market histories face different trade-offs.

A worker in her true calling faces a relatively straightforward separation decision. That is, she knows to be working in her best occupational fit, as such there is no longer any scope for learning. As a result, the separation decision at the individual match level is akin to that in the standard MP model, where the trade-off is between the expected present discounted value (PDV) of output produced in a continuing match and the value of non-market activity. Specifically, match surplus generated by the worker in her true calling is

$$S_M(x) = \max \{S_M^c(x), 0\}, \quad (7)$$

with

$$S_M^c(x) = y_G + x - z + \beta \left[(1 - \delta) \int S_M(x) dF(x) - \tau p(\theta) S_M(x_u) \right], \quad (8)$$

where $S_M^c(x)$ is the continuation surplus generated by a worker in her best occupational fit. For a worker in the learning phase, instead, the separation decision involves additional considerations. By destroying the match, the worker forgoes not only the expected PDV of output that would be otherwise produced in the match, but also the possibility to learn whether the current occupation is her true calling or not (“value of learning”). Specifically, match surplus during the learning phase for a worker of type $i = 1, \dots, M - 1$ is

$$S_{L,i}(x) = \max \{S_{L,i}^c(x), 0\}, \quad (9)$$

with

$$\begin{aligned} S_{L,i}^c(x) = & \bar{y}_i + x - z + \beta \left[(1 - \lambda)(1 - \delta) \int S_{L,i}(x) dF(x) - \tau p(\theta) S_{L,i}(x_u) \right] + \\ & + \underbrace{\beta \lambda \left[\pi_i \int S_M(x) dF(x) + (1 - \pi_i) \int S_{B,i}(x) dF(x) \right]}_{\text{value of learning}}, \end{aligned} \quad (10)$$

where $S_{L,i}^c(x)$ is the continuation surplus of a match of unknown quality, and the last term on the right-hand side of equation (10) is what we refer to as value of learning.

For a worker in a bad occupational match, instead, the separation decision has to take in account that by remaining in the current match, the worker forgoes the value of trying a new occupation that could eventually lead to the discovery of her true calling (“value of experimentation”). Specifically, match surplus for the worker in a bad occupation is

$$S_{B,i}(x) = \max \{S_{B,i}^c(x), 0\}, \quad (11)$$

with

$$S_{B,i}^c(x) = y_B + x - z + \beta (1 - \delta) \int S_{B,i}(x) dF(x) - \underbrace{\beta \tau p(\theta) S_{L,i+1}(x_u)}_{\text{value of experimentation}}, \quad (12)$$

where $S_{B,i}^c(x)$ is the continuation surplus of a bad match, and the last term on the right-hand side of equation (12) is what we refer to as value of experimentation.

In free-entry equilibrium, employers post job vacancies up to the point where the expected cost equals the expected benefit of opening and maintaining a job vacancy:

$$\frac{k}{q(\theta)} = \beta (1 - \tau) \left[\sum_{i=1}^{M-1} \phi_i S_{L,i}(x_u) + \phi_M S_M(x_u) \right], \quad (13)$$

with

$$\phi_i \equiv \frac{\mu_i u_i}{\sum_{i=1}^M \mu_i u_i}, \quad (14)$$

where μ_i is the total mass of workers of type i and u_i is unemployment rate for workers of type i . The aggregate unemployment rate is $u = \sum_{i=1}^M \mu_i u_i$ and so the unemployment shares ϕ_i ’s account for the composition of the unemployment pool. We assume that all new jobs are created at the upper bound of the match-specific productivity distribution, x_u , such that match surpluses are guaranteed to be positive for all workers’ types. Equation (13) pins down market tightness, θ , and so the aggregate job-finding rate, $p(\theta)$.

3.3 Endogenous job separation

We now turn to the worker's job separation decision. Since the continuation values of the match surplus in good, bad, and unknown occupational matches are monotonically increasing in match-specific productivity, x , job separation satisfies the reservation property. That is, for each occupational match and worker's type, there exists a unique cutoff value, x_d , such that all matches with $x \leq x_d$ are endogenously destroyed. Since employers and workers have the option to separate at no cost, a match remains active as long as its continuation value is above zero. Hence, at the *individual* match level, the job-separation rate is

$$\delta(x) = \begin{cases} \delta & \text{if } x > x_d \\ 1 & \text{if } x \leq x_d. \end{cases} \quad (15)$$

Since match surplus is increasing in the quality of the match, the cutoff value x_d varies across occupational matches: $x_d^M < x_d^{L,i} < x_d^{B,i}$, i.e., workers in the best occupational fit are less likely to separate than workers in unknown occupational matches during the learning phase, and in turn workers in the learning phase are less likely to separate than workers in bad occupational matches.

Specifically, for workers in good occupational matches, the expected job-separation rate is

$$\bar{\delta}_M \equiv \mathbb{E}_x [\delta_M(x)] = \delta + (1 - \delta)F(x_d^M), \quad (16)$$

where $F(x_d^M)$ is the mass of workers below the productivity cutoff x_d^M . For workers in bad and unknown occupational matches, instead, expected job-separation rates are, respectively,

$$\bar{\delta}_{B,i} \equiv \mathbb{E}_x [\delta_{B,i}(x)] = \delta + (1 - \delta)F(x_d^{B,i}), \quad (17)$$

and

$$\bar{\delta}_{L,i} \equiv \mathbb{E}_x [\delta_{L,i}(x)] = \delta + (1 - \delta)F(x_d^{L,i}), \quad (18)$$

with $\bar{\delta}_M < \bar{\delta}_{L,i} < \bar{\delta}_{B,i}$. Note that a worker of type $i < M$ can be either in a bad occupational match, with probability $\lambda(1 - \pi_i)$, or in an unknown occupational match, with probability $(1 - \lambda)$, such that type-specific, expected job-separation rates are

$$\bar{\delta}_i = \lambda(1 - \pi_i) \bar{\delta}_{B,i} + (1 - \lambda) \bar{\delta}_{L,i} \quad \text{for } i < M, \quad (19)$$

and $\bar{\delta}_i = \bar{\delta}_M$ for $i = M$.

3.4 Micro and macro effects of tax changes

We now turn to the propagation mechanism of tax shocks embodied in the model. We discuss first the effects of changes in taxes on type-specific unemployment rates (“micro effects”) and then those on the aggregate unemployment rate (“macro effects”).

Micro effects. In the stationary equilibrium, the unemployment rate for worker’s type i is

$$u_i = \frac{\bar{\delta}_i}{\bar{\delta}_i + p(\theta)}, \quad (20)$$

where $\bar{\delta}_i$ is the type-specific, expected job-separation rate, and $p(\theta)$ is the aggregate job-finding rate. Differentiation of (20) with respect to the tax rate, τ_W , yields

$$\frac{du_i}{d\tau_W} = \underbrace{(1 - u_i) u_i \times \frac{d \ln \bar{\delta}_i}{d\tau_W}}_{\text{separation-driven tax effect}} - \underbrace{(1 - u_i) u_i \epsilon_{p,\theta} \times \frac{d \ln \theta}{d\tau_W}}_{\text{hiring-driven tax effect}}, \quad (21)$$

where $\epsilon_{p,\theta}$ is the elasticity of the aggregate job-finding rate with respect to tightness. Equation (21) decomposes the effects of a tax change on type-specific unemployment rates: (i) changes in the tax rate affect job separation decisions at the individual match level, and so type-specific expected job-separation rates, as they change the value of the surplus generated by the match. This effect is captured by the semi-elasticity $d \ln \bar{\delta}_i / d\tau_W$ —*separation-driven tax effect*; and (ii) changes in the tax rate also affect vacancy posting and, as a result, market tightness, and so the aggregate job-finding rate, through the free-entry condition (13). This effect is captured by the semi-elasticity $d \ln \theta / d\tau_W$ —*hiring-driven tax effect*.

The relevant semi-elasticities for the separation-driven tax effect are

$$\frac{d \ln \bar{\delta}_M}{d\tau_W} = \frac{(1 - \delta) F(x_d^M)}{\bar{\delta}_M} \times \frac{f(x_d^M)}{F(x_d^M)} \times \frac{dx_d^M}{d\tau_W} \quad \text{for } i = M, \quad (22)$$

$$\frac{d \ln \bar{\delta}_{B,i}}{d\tau_W} = \frac{(1-\delta) F(x_d^{B,i})}{\bar{\delta}_{B,i}} \times \lambda (1-\pi_i) \frac{f(x_d^{B,i})}{F(x_d^{B,i})} \times \frac{dx_d^{B,i}}{d\tau_W} \quad \text{for } i < M, \quad (23)$$

and

$$\frac{d \ln \bar{\delta}_{L,i}}{d\tau_W} = \frac{(1-\delta) F(x_d^{L,i})}{\bar{\delta}_{L,i}} \times (1-\lambda) \frac{f(x_d^{L,i})}{F(x_d^{L,i})} \times \frac{dx_d^{L,i}}{d\tau_W} \quad \text{for } i < M, \quad (24)$$

where the first terms on the right-hand side of equations (22), (23), and (24) are approximately equal to 1 for “small” values of the exogenous job-separation rate, δ .

Equations (22)-(24) identify two channels through which taxes affect type-specific, mean job-separation rates: (i) a change in the tax rate causes a change in the productivity cutoff for endogenous job separations. This effect is captured by the term $dx_d/d\tau_W$ —*cutoff-driven tax effect*; and (ii) the effect of any change in the productivity cutoff depends on a measure of homogeneity of the workforce at the margin (akin to a hazard rate). This effect is captured by the term $f(x_d)/F(x_d)$ —*hazard-driven tax effect*—that is the ratio of the density of workers at the margin to the mass of workers on the left of the productivity cutoff.¹¹

The relevant semi-elasticity for the hiring-driven tax effect is

$$\frac{d \ln \theta}{d\tau_W} = \frac{1}{(1-\alpha)\bar{S}(x_u)} \times \left[\sum_{i=1}^{M-1} \phi_i \frac{dS_{L,i}(x_u)}{d\tau_W} + S_{L,i}(x_u) \frac{d\phi_i}{d\tau_W} + \phi_M \frac{dS_M(x_u)}{d\tau_W} + S_M(x_u) \frac{d\phi_M}{d\tau_W} \right],$$

where $\bar{S}(x_u) \equiv \sum_{i=1}^{M-1} \phi_i S_{L,i}(x_u) + \phi_M S_M(x_u)$ is a weighted average of match surpluses across all worker’s types. Note that using $\sum_{i=1}^M \phi_i = 1$, and

$$\frac{dS_{L,i}(x_u)}{d\tau_W} = \frac{dS_M(x_u)}{d\tau_W} = -\frac{dz}{d\tau_W} \quad \text{for all } i < M,$$

the equation becomes

$$\frac{d \ln \theta}{d\tau_W} = \frac{1}{(1-\alpha)\bar{S}(x_u)} \times \left[-\frac{dz}{d\tau_W} + \sum_{i=1}^{M-1} S_{L,i}(x_u) \frac{d\phi_i}{d\tau_W} + S_M(x_u) \frac{d\phi_M}{d\tau_W} \right], \quad (25)$$

which identifies two effects of tax changes on market tightness: (i) a direct effect on match surpluses through worker’s net productivity, captured by the term $dz/d\tau_W$ —*net-productivity-driven tax effect*; and (ii) a change in the tax rate leads to changes in unemployment shares

¹¹Note that under our assumption of uniformly distributed match-specific shocks, the density of workers at the margin, $f(x_d)$, is in fact independent of the specific value of the productivity cutoff, x_d .

and so it affects the composition of the unemployment pool—*composition-driven tax effect*.

Macro effects. To investigate the aggregate effects of tax changes, we need first to discuss aggregation. Specifically, to calculate the aggregate unemployment rate, we need the shares of each worker’s type in the stationary equilibrium of the model. At each point in time, the worker’s type distribution evolves stochastically according to an endogenous Markov chain. As an illustrative example, below we show the type transition matrix for $M = 3$,

$$[\mu_{i,i'}] = \begin{bmatrix} 1 - \sigma + \sigma(1 - \lambda) + \sigma\lambda(1 - \pi_1)(1 - \bar{\delta}_{B,1}) & \sigma\lambda(1 - \pi_1)\bar{\delta}_{B,1} & \sigma\lambda\pi_1 \\ 1 - \sigma & \sigma(1 - \lambda) & \sigma\lambda \\ 1 - \sigma & 0 & \sigma \end{bmatrix},$$

where $\mu_{i,i'} = \text{prob}\{i_{t+1} = i' | i_t = i\}$ is the probability that a worker of type i becomes of type i' next period. The transition matrix $[\mu_{i,i'}]$ produces a non-degenerate, stationary distribution of worker’s types μ , whose i -th element, μ_i , corresponds to the mass of workers of type i in the stationary equilibrium. Hence, the aggregate unemployment rate is

$$u = \sum_{i=1}^M \mu_i u_i. \quad (26)$$

Log-differentiation of (26) with respect to τ_W , rearranging terms, and using $\phi_i \equiv \mu_i u_i / u$, yields

$$\frac{d \ln u}{d \tau_W} = \sum_{i=1}^M \phi_i \left(\frac{d \ln \mu_i}{d \tau_W} + \frac{d \ln u_i}{d \tau_W} \right), \quad (27)$$

which decomposes the aggregate unemployment semi-elasticity to tax changes into two components: (i) the semi-elasticity of the shares of each worker’s type; and (ii) the semi-elasticity of type-specific unemployment rates, that is, the micro effect in equation (21) (all weighted by type-specific unemployment shares). Equation (27) suggests that the aggregate unemployment semi-elasticity in equation (27) is a complicated function of a large number of structural parameters.

4 Quantitative Analysis

In this section, we study the quantitative properties of the theory. The model is calibrated to post-war U.S. data for 1950-2006. We show that the theory quantitatively explains the observed life-cycle profile of average unemployment rates and that it replicates the *estimated* semi-elasticity of the aggregate unemployment rate to tax shocks.

4.1 Model calibration

We now turn to the parameterization of the model. Many features of the model are standard to the search-and-matching literature, as such I aim to maintain comparability with previous work whenever possible.

The discount rate β is set to match an annual interest rate of 4 percent as in [Kydland and Prescott \(1982\)](#). A large literature that directly estimates the aggregate matching function, provides estimates for the parameter α , that also corresponds to the elasticity of hiring with respect to job vacancies. For instance, [Petrongolo and Pissarides \(2001\)](#) establish a plausible range of 0.3-0.5, whereas [Brügemann \(2008\)](#) obtains a refined range of 0.37-0.46. We then set α to 0.4 at the mid point of these ranges. As standard in the literature, we set the worker's Nash-bargaining weight τ to $1 - \alpha$, such that the [Hosios \(1990\)](#)'s condition is met and the decentralized equilibrium, without taxes, is constrained efficient. Based on [Shimer \(2005\)](#), we target a tightness ratio of 1, which requires setting the cost of posting a job vacancy k to 0.2192. The matching function scale parameter ξ is then set to 0.453 to match the average job-finding rate in the data for 1950-2006. As for the labor income tax, we set τ_W to 0.2972, that is the average of AMTR for 1950-2006 as in [Barro and Redlick \(2011\)](#).

The parameters governing flow value of unemployment, b , exogenous job-separation rate, δ , and upper bound of the distribution for match-specific shocks, x_u , are calibrated to match an average replacement ratio of 40 percent as in [Shimer \(2005\)](#), the average job-separation rate of 2.7 percent and the average unemployment rate of 4.9 percent for workers of 20-64 years old. This requires setting b to 0.33, δ to 0.01, and x_u to 0.129. Moreover, we set the survival probability σ to 0.9986 for the model to reproduce the average age of the labor force for 1950-2006, which is 39.3 in the data and turns out to be 39.5 in the model.

Finally, we calibrate the remaining parameters of the model, i.e., y_G , y_B , and λ , to match the empirical life-cycle profile of earnings, as estimated by [Murphy and Welch \(1990\)](#) and many others. As such, we associate returns to occupational learning with returns to labor

market experience. Specifically, we target two key properties of the return to experience: (i) the maximal wage gain for a typical worker over the life cycle represents an approximate doubling of earnings; and (ii) such doubling occurs after roughly 25 years of experience in the labor market. After normalizing y_G to 1, matching those targets requires setting y_B to 0.485 and λ to 18^{-1} . These parameter values imply that a typical worker is, on average, nearly twice as productive in her best occupational fit as in the bad occupational match, and that it takes, on average, one and a half years in a match to learn the quality of such a match.

4.2 Life-cycle unemployment and aggregate unemployment semi-elasticity

We now turn to discuss the model's ability to replicate the observed life-cycle profile of unemployment rates and the estimated aggregate unemployment semi-elasticity to tax cuts.

In Figure 7, panel A shows that the model is able to reproduce the (average of) age-specific unemployment rates observed in the data for the period 1950-2006. In the model, as in the data, unemployment rates decrease monotonically with age. The average unemployment rate of 9 percent for the 20-24 years old is approximately 2.5 times larger than the 3.6 percent for the 55-64 years old. Moreover, in the model, as in the data, the relationship between unemployment rates and age is convex: average unemployment rates decrease sharply from the 20-24 to 25-34 years old and remain essentially flat for the 35-64 years old.

In Figure 7, panel B shows that the model produces a marked life-cycle profile of job-separation rates: the young experience higher job-separation rates than prime-age workers. Note that, in the model, the job-finding rate does not vary across age groups (see panel C, Figure 7). Thus, the key driving force of the large age-specific differences in unemployment rates generated by the model is the life-cycle profile of job-separation rates, which is in fact consistent with previous studies on unemployment over the life cycle (see [Menzio et al., 2012](#); [Gervais et al., 2014](#)).

In Figure 7, panel D shows that the model generates a concave life-cycle profile of match surpluses. During the youth, workers switch jobs to learn the occupation in which they are most productive. Hence, young workers are most likely to be in bad occupational matches that generate small surplus. In contrast, prime-age workers are most likely settled in their best occupational fit, and so generating large surplus.

In Table 2, panel B shows that the calibrated model matches the estimated semi-elasticity

Figure 7: Age-specific Unemployment Rates

Notes: Panel A shows age-specific unemployment rates in the model (full line with diamonds) and in the data for 1950-2006 (full line with circles). Panel B and C show age-specific job-separation and job-finding rates, respectively. Panel D shows age-specific match surplus.

Table 2: Life-cycle Unemployment and Semi-elasticities

	Data	Model
A. Unemployment rates (percent)		
$\bar{u}_{20-24}$	9.01	9.44
$\bar{u}_{25-34}$	5.29	5.60
$\bar{u}_{35-44}$	4.04	3.60
$\bar{u}_{45-54}$	3.64	3.33
$\bar{u}_{55-64}$	3.62	3.32
B. Semi-elasticities		
$d \ln u / d\tau_W$	6.67	6.67
$d \ln s / d\tau_W$	—	5.04
$d \ln f / d\tau_W$	—	2.10

Notes: In panel A, all reported statistics refer to the sample period 1950-2006; $\bar{u}_a$ indicates the sample average of the unemployment rate for the age group $a \in \{20-24, 25-34, 35-44, 45-54, 55-64\}$. In panel B, the unemployment semi-elasticity in the data is that implied by the estimates based on the counterfactual unemployment rate series for the age group 20-64 with the age-specific labor force shares fixed at their sample averages for 1950-2006, see equation (6) in Section 2.

of the aggregate U.S. unemployment rate to tax shocks. As such, in Section 5, we use the calibrated model to reconstruct the U.S. history of unemployment semi-elasticities as we vary the average age of the labor force consistently with the data.

5 Baby Boom and Unemployment Semi-elasticities

In this section, we study how the aging of the U.S. labor force impact the propagation of tax shocks. The empirical results in Shimer (1999) indicate that the age composition of the labor force has a causal impact on the *level* of the aggregate U.S. unemployment rate. That is, the entry of the baby boom into the labor force in the late-1970s and their subsequent aging accounts for a substantial fraction of the rise and fall in the (average) unemployment rates observed in past 50 years. In addition, Jaimovich and Siu (2009) argue that the age composition of the labor force has a causal impact on the *volatility* of the U.S. business cycle. Since young workers feature less volatile hours worked than prime-age workers, the aging of the labor force accounts for a significant fraction of the decrease in business cycle volatility since the mid-1980s over the so-called Great Moderation.

We next show that the demographic change experienced by the U.S. labor market is relevant for the propagation mechanism of tax shocks. To this aim, we proceed in two steps. First, we fit the model to the age-specific labor force shares in different years over the period 1950-2015, which requires recalibrating the survival probability σ to match the average age of the labor force in those years. We keep all the other parameters fixed at their baseline values. Second, through the lens of the calibrated model, we reconstruct the U.S. history of aggregate unemployment semi-elasticities.

In Table 3, we report the semi-elasticities for the job-separation and job-finding rates, and the aggregate unemployment rate implied by the model from 1950 to 2015 at 10-year intervals. Over this period, the United States have experienced dramatic changes in the age composition of the labor force. For instance, the average age of the U.S. labor force was 37.79 in 1980 and raised as high as 41.21 in 2015. The results indicate a tight link between the semi-elasticity of the aggregate unemployment rate to tax shocks and the average age of the labor force. Specifically, the aggregate unemployment semi-elasticity in 2015 is estimated to be 4.55 as opposed to 7.50 in 1980. That is, the calibrated model implies that the aging of the Baby Boom decreases the aggregate unemployment semi-elasticity by approximately 50 percent.

Table 3: Baby Boom and Unemployment Semi-elasticities

Year	Average age labor force	Semi-elasticity: $d \ln u / d\tau_W$	Semi-elasticity: $d \ln s / d\tau_W$	Semi-elasticity: $d \ln f / d\tau_W$
1950	39.37	6.64	5.01	2.09
1960	40.72	5.12	3.97	1.46
1970	40.05	5.77	4.42	1.71
1980	37.79	7.50	5.48	2.65
1990	37.91	7.20	5.22	2.56
2000	39.66	6.36	4.83	1.96
2010	41.10	4.65	3.66	1.25
2015	41.21	4.55	3.59	1.20

Notes: The average age of the labor force is calculated as $\sum_a \left(\frac{a+\bar{a}}{2} \right) \text{LFS}_a$, where $\underline{a}$ and $\bar{a}$ are the lower and upper age bounds of the age group a , respectively, and LFS_a is the share in the labor force for the 20-64 age group of the age group $a \in \{20-24, 25-34, 35-44, 45-54, 55-64\}$.

6 Conclusion

In this paper, we investigate the consequences of demographic change for the propagation of tax shocks. After isolating exogenous variation in average marginal tax rates in SVARs using a narrative identification approach, we document that the responsiveness of unemployment rates to tax changes largely varies across age groups: the unemployment rate response of the young is nearly twice as large as that of prime-age workers. We then present a theory that quantitatively accounts for the estimated semi-elasticity of the aggregate unemployment rate to tax shocks. Through the lens of the calibrated model, we reconstruct the U.S. history of aggregate unemployment semi-elasticities from 1950 to 2015. The results indicate that changes in the age composition of the U.S. labor force constitute an important propagation mechanism of tax shocks.

Appendix

A Data Sources

Monthly data for the U.S. labor force, unemployment rates, labor force participation rates, employment and unemployment levels (number in thousands) for the overall economy and by age groups are produced by the Bureau of Labor Statistics (BLS) and publicly available at the Current Population Survey (CPS) home page at <http://www.bls.gov/cps/>. Monthly job-separation and job-finding rates are constructed based on Shimer (2012) and publicly available at Robert Shimer's website at <https://sites.google.com/site/robertshimer/research/flows>. Data on average marginal income tax rates has been tabulated by Robert Barro and Charles Redlick in Barro and Redlick (2011). Data for personal income tax base, government purchases, and federal debt (all in real per capita terms) are from Mertens and Ravn (2013) dataset and publicly available at Karel Mertens's website at <https://mertens.economics.cornell.edu/research.htm>.

B Additional Tables and Figures

Figure B.1: Output Response to a Tax Cut

Notes: The figure shows the response to a 1 percentage point cut in the average marginal personal income tax rate. Full lines with circles are point estimates; dashed lines are 95 percent confidence bands.

Figure B.2: Age-specific Participation Rate Response to a Tax Cut

Notes: The figure shows the response to a 1 percentage point cut in the average marginal personal income tax rate. Full lines with circles are point estimates; dashed lines are 95 percent confidence bands.

References

- E. Anderson, A. Inoue, and B. Rossi. Heterogeneous Consumers and Fiscal Policy Shocks. *Unpublished manuscript*, 2011.
- A.J. Auerbach and Y. Gorodnichenko. Measuring the Output Responses to Fiscal Policy. *American Economic Journal: Economic Policy*, 4(2):1–27, 2012.
- R. Barnichon and C. Matthes. Stimulus versus Austerity: The Asymmetric Government Spending Multiplier. *Unpublished manuscript*, 2015.
- R.J. Barro and C.J. Redlick. Macroeconomic Effects from Government Purchases and Taxes. *Quarterly Journal of Economics*, 126(1):51–102, 2011.
- B. Brügemann. What Elasticity of the Matching Function is Consistent with US Aggregate Labor Market Data? *Unpublished manuscript*, 2008.
- C. Burnside, M. Eichenbaum, and J.D.M. Fisher. Fiscal Shocks and their Consequences. *Journal of Economic theory*, 115(1):89–117, 2004.
- C.F. Christ. A Simple Macroeconomic Model with a Government Budget Restraint. *Journal of Political Economy*, 76(1):53–67, 1968.
- J. Cloyne. Discretionary Tax Changes and the Macroeconomy: New Narrative Evidence from the United Kingdom. *American Economic Review*, 103(4):1507–1528, 2013.
- S. Dyrda, G. Kaplan, and J.V. Ríos-Rull. Business Cycles and Household Formation: The Micro vs the Macro Labor Elasticity. *NBER Working Paper Series*, 17880, 2012.
- M.W.L. Elsby, R. Michaels, and G. Solon. The Ins and Outs of Cyclical Unemployment. *American Economic Journal: Macroeconomics*, 1(1):84–110, 2009.
- S. Fujita and G. Ramey. The Cyclicalities of Separation and Job Finding Rates. *International Economic Review*, 50(2):415–430, 2009.
- M. Gervais, N. Jaimovich, H.E. Siu, and Y. Yedid-Levi. What Should I Be When I Grow Up? Occupations and Unemployment over the Life Cycle. *NBER Working Paper Series*, 20628, 2014.

- P. Gomme, R. Rogerson, P. Rupert, and R. Wright. The Business Cycle and the Life Cycle. *NBER Macroeconomics Annual 2004*, 19:415–461, 2005.
- R.E. Hall. Employment Efficiency and Sticky Wages: Evidence from Flows in the Labor Market. *Review of Economics and Statistics*, 87(3):397–407, 2005.
- G.D. Hansen and S. İmrohoroglu. Business Cycle Fluctuations and the Life Cycle: How Important is On-the-job Skill Accumulation? *Journal of Economic Theory*, 144(6):2293–2309, 2009.
- A.J. Hosios. On the Efficiency of Matching and Related Models of Search and Unemployment. *Review of Economic Studies*, 57(2):279–298, 1990.
- N. Jaimovich and H.E. Siu. The Young, the Old, and the Restless: Demographics and Business Cycle Volatility. *American Economic Review*, 99(3):804–826, 2009.
- N. Jaimovich, S. Pruitt, and H.E. Siu. The Demand for Youth: Explaining Age Differences in the Volatility of Hours. *American Economic Review*, 103(7):3022–3044, 2013.
- F.E. Kydland and E.C. Prescott. Time to Build and Aggregate Fluctuations. *Econometrica*, 50(6):1345–1370, 1982.
- G. Menzio, I.A. Telyukova, and L. Visschers. Directed Search over the Life Cycle. *NBER Working Paper Series*, 17746, 2012.
- K. Mertens. Marginal Tax Rates and Income: New Time Series Evidence. *NBER Working Paper Series*, 19171, 2013.
- K. Mertens and M.O. Ravn. Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated US Tax Policy Shocks. *American Economic Journal: Economic Policy*, 4(2):145–181, 2012.
- K. Mertens and M.O. Ravn. The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States. *American Economic Review*, 103(4):1212–1247, 2013.
- K. Mertens and M.O. Ravn. A Reconciliation of SVAR and Narrative Estimates of Tax Multipliers. *Journal of Monetary Economics*, 68(S):S1–S19, 2014.
- T. Monacelli, R. Perotti, and A. Trigari. Unemployment Fiscal Multipliers. *Journal of Monetary Economics*, 57(5):531–553, 2010.

- D.T. Mortensen and C.A. Pissarides. Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies*, 61(3):397–415, 1994.
- K.M. Murphy and F. Welch. Empirical Age-Earnings Profiles. *Journal of Labor Economics*, 8(2):202–229, 1990.
- B. Petrongolo and C.A. Pissarides. Looking into the Black Box: A Survey of the Matching Function. *Journal of Economic Literature*, 39(2):390–431, 2001.
- C.A. Pissarides. Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages. *American Economic Review*, 75(4):676–690, 1985.
- C.A. Pissarides. The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer? *Econometrica*, 77(5):1339–1369, 2009.
- M.J. Pries. Persistence of Employment Fluctuations: A Model of Recurring Job Loss. *Review of Economic Studies*, 71(246):193, 2004.
- V.A. Ramey. Identifying Government Spending Shocks: It’s all in the Timing. *Quarterly Journal of Economics*, 126(1):1–50, 2011a.
- V.A. Ramey. Can Government Purchases Stimulate the Economy? *Journal of Economic Literature*, 49(3):673–685, 2011b.
- V.A. Ramey and M.D. Shapiro. Costly Capital Reallocation and the Effects of Government Spending. *Carnegie-Rochester Conference Series on Public Policy*, 48:145–194, 1998.
- V.A. Ramey and S. Zubairy. Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data. *NBER Working Paper Series*, 20719, 2014.
- C.D. Romer and D.H. Romer. The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks. *American Economic Review*, 100(3):763–801, 2010.
- R. Shimer. Why is the U.S. Unemployment Rate So Much Lower? *NBER Macroeconomics Annual 1998*, 13:11–74, 1999.
- R. Shimer. The Cyclical Behavior of Equilibrium Unemployment and Vacancies. *American Economic Review*, 95(1):25–49, 2005.
- R. Shimer. Reassessing the Ins and Outs of Unemployment. *Review of Economic Dynamics*, 15(2):127–148, 2012.

C.A. Sims. Econometric Implications of the Government Budget Constraint. *Journal of Econometrics*, 83(1):9–19, 1998.